

Lower Bounds for Monotone Counting Circuits*

Stasys Jukna[†]

February 9, 2015

Abstract

A $\{+, \times\}$ -circuit *counts* a given multivariate polynomial f , if its values on 0-1 inputs are the same as those of f ; on other inputs the circuit may output arbitrary values. Such a circuit counts the number of monomials of f evaluated to 1 by a given 0-1 input vector (with multiplicities given by their coefficients). A circuit *decides* f if it has the same 0-1 roots as f . We first show that some multilinear polynomials can be exponentially easier to count than to compute them, and can be exponentially easier to decide than to count them. Then we give general lower bounds on the size of counting circuits.

Keywords: arithmetic circuits, boolean circuits, counting complexity, lower bounds

1 Introduction

In this paper we consider computational complexity of multivariate polynomials with nonnegative integer coefficients:

$$f(x_1, \dots, x_n) = \sum_{e \in \mathbb{N}^n} c_e \prod_{i=1}^n x_i^{e_i}, \quad (1)$$

where $c_e \in \mathbb{N} = \{0, 1, 2, \dots\}$, and $x_i^0 = 1$. Products $\prod_{i=1}^n x_i^{e_i}$ are *monomials* of f ; we will often omit monomials whose coefficients c_e are zero. The polynomial is *multilinear*, if $c_e = 0$ for all $e \notin \{0, 1\}^n$, and is *homogeneous* of degree d , if $e_1 + \dots + e_n = d$ for all e with $c_e \neq 0$.

A standard model of compact representation of such polynomials (with nonnegative coefficients) is that of monotone arithmetic circuits, i.e. of $\{+, \times\}$ -circuits. Such a circuit is a directed acyclic graph with three types of nodes: input, addition (+), and multiplication (\times). Input nodes have fanin zero, and correspond to variables x_1, \dots, x_n . All other nodes have fanin two, and are called *gates*. The *size* of a circuit is the number of gates in it.

Every $\{+, \times\}$ -circuit syntactically *produces* a unique monotone polynomial F with nonnegative integer coefficients in a natural way: the polynomial produced at an input gate x_i consists of a single monomial x_i , and the polynomial produced at a sum (product) gate is the sum (product) of polynomials produced at its inputs; we use distributivity to write a product of polynomials as a sum of monomials. The polynomial F produced by the circuit itself is the polynomial produced at its output gate. Given a polynomial $f(x_1, \dots, x_n)$, we say that the circuit

*Research supported by the DFG grant SCHN 503/6-1.

[†]Institute of Computer Science, Goethe University, Frankfurt am Main, Germany. Affiliated with Institute of Mathematics and Informatics, Vilnius University, Vilnius, Lithuania. Email: jukna@thi.informatik.uni-frankfurt.de




Figure 1: A circuit of size 5 counting the polynomial $f = 2xyz + 2xy + 2xz + 2yz$, and deciding the polynomial $g = xy + xz + yz$. The circuit itself produces the polynomial $F = (x + y)(y + z)(x + z) = 2xyz + x^2y + xy^2 + x^2z + xz^2 + y^2z + yz^2$. Gate u is the output gate.

- *computes* f , if $F(a) = f(a)$ holds for all $a \in \mathbb{N}^n$, where $\mathbb{N} = \{0, 1, 2, \dots\}$;
- *counts* f , if $F(a) = f(a)$ holds for all $a \in \{0, 1\}^n$;
- *decides* f , if $F(a) = 0$ exactly when $f(a) = 0$ holds for all $a \in \{0, 1\}^n$.

In this paper we are mainly interested in $\{+, \times\}$ -circuits *counting* a given polynomial f . Such a circuit needs only to correctly compute the restriction $f : \{0, 1\}^n \rightarrow \mathbb{N}$ of f on 0-1 inputs. Note that, if the polynomial f is monic (has no coefficients > 1) then, on every 0-1 input $a \in \{0, 1\}^n$, $f(a)$ is the number of monomials of f satisfied by (evaluated to 1 on) a . For example, in the case of the permanent polynomial

$$\text{Per}_n(x) = \sum_h \prod_{i=1}^n x_{i,h(i)}$$

with the summation over all permutations h of $[n] = \{1, \dots, n\}$, the value $\text{Per}_n(a)$ is the number of perfect matchings in the bipartite $n \times n$ graph G_a specified by input $a \in \{0, 1\}^{n \times n}$; nodes i and j are adjacent in G_a if and only if $a_{ij} = 1$. Thus, a circuit counting Per outputs the number of perfect matchings in G_a , whereas a circuit deciding this polynomial merely tells us whether G_a contains a perfect matching.

Remark 1. Let us stress that we only consider *monotone* arithmetic circuits. The reason is that counting $\{+, -, \times\}$ -circuits are already omnipotent: they are as powerful as boolean $\{\vee, \wedge, \neg\}$ -circuits. This is because each boolean operation can be simulated over $\{0, 1\}$: $x \wedge y$ by xy , $\neg x$ by $1 - x$, and $x \vee y$ by $x + y - xy$.

If a $\{+, \times\}$ -circuit computes, counts or only decides a given polynomial f , what can then be said about the structure of the produced by the circuit polynomial F ?

If the circuit *computes* f , then $F = f$ must hold, that is, then the produced polynomial F and the target polynomial f must coincide as formal expressions, i.e. as sums of monomials (see, e.g. Claim 10 below for simple a proof). In particular, then $\text{mon}(F) = \text{mon}(f)$ must also hold, where

- $\text{mon}(f)$ is the set of monomials appearing in f with nonzero coefficients.

This ensures that no “invalid” monomials can be formed during the computation, and severely limits the power of such circuits. In particular, if the target polynomial f is *multilinear* (no variable has degree larger than 1, then the circuit itself must be multilinear: the polynomials produced at inputs of each product gate must depend on *disjoint* sets of variables. This limitation is essentially exploited in all lower bounds for monotone arithmetic circuits, including [17, 19, 10, 27, 22, 6, 24, 7].

In counting circuits, $\text{mon}(F) = \text{mon}(f)$ needs not to hold, due to the multiplicative idempotence axiom $x^2 = x$ valid on 0-1 inputs. That is, here exponents (and hence, degrees of monomials) do not mater (see Fig. 1). Still, it can be shown (see Lemma 5 below) that here we have a weaker, but still strong enough property $\text{sup}(F) = \text{sup}(f)$, where

- $\text{sup}(f)$ is the *support* of f , that is, the family of sets of variables of monomials in $\text{mon}(f)$.

In deciding circuits, even $\text{sup}(F) = \text{sup}(f)$ needs not to hold, due to the additional absorption axiom $x + xy = x$. In such circuits, we only have a weak property $\text{min}(F) = \text{min}(f)$, where

- $\text{min}(f) \subseteq \text{sup}(f)$ is the family of all members of $\text{sup}(f)$ which are minimal in the sense that they do not contain any other members of $\text{sup}(f)$; hence, $\text{min}(f)$ forms an antichain.

Deciding $\{+, \times\}$ -circuits are actually monotone *boolean* circuits, and we have the following relations concerning the minimum circuit size for every given polynomial (we will prove that both gaps can be exponential):

$$\text{Deciding} \leq \text{Counting} \leq \text{Computing}.$$

To prove lower bounds for deciding, and hence, also for counting $\{+, \times\}$ -circuits, one can use lower-bounds arguments for monotone boolean circuits (see, e.g. [12, Chapt. 9] and the literature cited herein), but these are not easy to apply. The reason here lies in a “dual character” of these arguments: in order to obtain a large lower bound on the decision complexity of a given polynomial f , not only the set of monomials of the polynomial f itself but also that of the “dual” polynomial f^* must have some good structural properties (see the discussion before Theorem 8 below).

On the other hand, due to the limitations we mentioned above, lower bounds for $\{+, \times\}$ -circuits *computing* a given polynomial are relatively easy to obtain, because here we have a full knowledge about the polynomial which a circuit must produce. In particular, there is then no need to consider dual polynomials. *Counting* $\{+, \times\}$ -circuits allow more freedom, because they can use $x^2 = x$. In this case we only know the structure of the support of the produced polynomial, but not about its monomials. So, it is natural to ask whether known lower bounds for exactly computing $\{+, \times\}$ -circuits can be extended to counting circuits?

That they sometimes *can* be extended was demonstrated by Sengupta and Venkateswaran in [18], where they show that an exponential lower bound of Jerrum and Snir [10] for $\{+, \times\}$ -circuits computing the permanent polynomial Per can be adopted to yield the same lower bound for circuits only counting this polynomial. Still, at least three questions remained open:

1. Can counting circuits be substantially smaller than computing circuits?
2. Can deciding circuits be substantially smaller than counting circuits?
3. Can lower-bounds *arguments* for computing $\{+, \times\}$ -circuits, not just bounds for specific polynomials (like the permanent polynomial), be extended to $\{+, \times\}$ -counting circuits?

In this paper, we answer these questions affirmatively.

2 Results

For a polynomial f , let $C(f)$ denote the minimum size of a $\{+, \times\}$ -circuit *computing* f , $C_{0/1}(f)$ the minimum size of such a circuit *counting* f , and $D(f)$ the minimum size of a $\{+, \times\}$ -circuit *deciding* f . Note that, for every polynomial f , we have that

$$D(f) \leq C_{0/1}(f) \leq C(f).$$

We first show that the gaps $C(f)/C_{0/1}(f)$ as well as $C_{0/1}(f)/D(f)$ can be exponential. When doing this, we will use known lower bound for the permanent polynomial.

Theorem 1 ([10, 18]). *If $f = \text{Per}_n$, then $C_{0/1}(f) \geq n2^{n-1}$.*

This lower bound on $C(f)$ was proved by Jerrum and Snir [10], and was extended to $C_{0/1}(f)$ by Sengupta and Venkateswaran [18] (see also Corollary 1 below for a short proof of a weaker $2^{\Omega(n)}$ lower bound).

We will also use the (simple) fact that it is not harder to compute the so-called “lower” and “higher” envelopes of polynomial than to compute the polynomial itself. The *lower envelope* of a polynomial f is a homogeneous polynomial f_{le} consisting of the monomials of f of smallest degree. The *higher envelope* f_{he} is defined by taking monomials of largest degree. (As usually, the degree of a monomial is the sum of exponents of its variables, and a polynomial is *homogeneous*, if all its monomials have the same degree.) As observed already by Jerrum and Snir [10], every $\{+, \times\}$ -circuit producing a polynomial f can be easily transformed into a circuit producing f_{le} or f_{he} by just discarding (if necessary) some of the sum-gates. Hence, we always have

$$C(f) \geq \max \{C(f_{\text{he}}), C(f_{\text{le}})\} . \quad (2)$$

2.1 Gaps

To show that the gap $C(f)/C_{0/1}(f)$ can be exponential, we will show a stronger fact that both gaps $C_{0/1}(f_{\text{he}})/C_{0/1}(f)$ and $C_{0/1}(f_{\text{le}})/C_{0/1}(f)$ can be exponential. Recall that, by (2), no such gap is possible for *computing* $\{+, \times\}$ -circuits.

Theorem 2. *There are multilinear polynomials f and g of n variables such that $C_{0/1}(f) = O(n)$ and $C_{0/1}(g) = O(n^{3/2})$, but both $C_{0/1}(f_{\text{he}})$ and $C_{0/1}(g_{\text{le}})$ are $2^{\Omega(\sqrt{n})}$.*

Remark 2. Together with (2), the theorem implies that the gap $C(f)/C_{0/1}(f)$ between the sizes of $\{+, \times\}$ -circuits computing and counting f can be exponential. Important in this result is that the gap is obtained for *multilinear* polynomials: this shows that, under the presence of multiplicative idempotence $x^2 = x$, non-multilinear circuits counting multilinear polynomials can be much more efficient. In this connection, let us mention that without this restriction (to multilinear polynomials) a non-trivial gap follows from the classical lower bound $\Omega(n \log d)$ of Strassen [23], and Baur and Strassen [2] on the size of arithmetic (not necessarily monotone) circuits computing the polynomial $f = x_1^d + x_2^d + \dots + x_n^d$, which can be trivially counted by a $\{+, \times\}$ -circuit $F = x_1 + x_2 + \dots + x_n$ of size only $n - 1$. But this example merely says that, under the presence of multiplicative idempotence $x^2 = x$, rising to powers is redundant.

To show that the gap $C_{0/1}(f)/D(f)$ can also be exponential, it is enough to take any polynomial $g(x_1, \dots, x_n)$ such that $C_{0/1}(g)$ is exponential, and consider the polynomial $f = g + h$ where $h = \sum_{i=1}^n x_i$. If $g(0, \dots, 0) = 0$ then, on every 0-1 input a , we have that $f(a) = 0$ if and only if $h(a) = 0$. So, f has a small decision complexity: $D(f) \leq D(h) \leq n$. So, if the counting complexity $C_{0/1}(f)$ of the extended polynomial f remains exponential, then the gap $C_{0/1}(f)/D(f)$ is exponential. In particular, one can establish such a gap by using the permanent polynomial $g = \text{Per}$ (the only small “technicality” here is to show that the counting complexity of f remains large).

Theorem 3. *If $f = \text{Per}_n + \sum_{i,j=1}^n x_{ij}$, then $D(f) \leq n^2$ but $C_{0/1}(f) = 2^{\Omega(n)}$.*

The polynomial used in this theorem is somewhat artificial. Actually, one can establish an exponential gap using a more natural (and important) s - t path polynomial Path_n . This polynomial has one variable $x_{i,j}$ for each edge of a complete undirected graph on $n+2$ nodes $\{s, 1, 2, \dots, n, t\}$. Each monomial of f corresponds to a simple directed path from node s to node t :

$$\text{Path}_n(x) = x_{s,t} + \sum_{l=1}^n \sum_{\substack{i_1, \dots, i_l \\ \text{distinct}}} x_{s,i_1} x_{i_1,i_2} \cdots x_{i_{l-1},i_l} x_{i_l,t}.$$

On a 0-1 input a , $\text{Path}_n(a)$ gives the number of s - t paths in the graph specified by a . Jerrum and Snir [10] have shown that every $\{+, \times\}$ -circuit *computing* $f = \text{Path}_n$ must have exponential size, i.e. that $C(f) = 2^{\Omega(n)}$. We show that even $\{+, \times\}$ -circuits *counting* Path must have exponential size.

Theorem 4. *If $f = \text{Path}_n$, then $D(f) = O(n^3)$, but $C_{0/1}(f) \geq 2^{n^{\Omega(1)}}$.*

2.2 Lower bounds

Recall that, if a $\{+, \times\}$ -circuit *computes* a given polynomial f , then the produced by the circuit polynomial F must just coincide with f (as formal expressions). In counting and deciding circuits we only have weaker conditions on F .

By the *linearization* of a polynomial f we will mean a multilinear polynomial \bar{f} obtained from f by removing all (nonzero) exponents from all monomials of f . For example, the linearization of $f = 2xy^2 + 3x^4y^2 + 6y^2z$ is $\bar{f} = 5xy + 6yz$. It is clear that $\bar{f}(a) = f(a)$ holds for all $a \in \{0, 1\}^n$.

Lemma 5. *If a $\{+, \times\}$ -circuit producing a polynomial F counts f , then $\bar{F} = \bar{f}$, and hence, also $\sup(F) = \sup(f)$. A $\{+, \times\}$ -circuit decides f if and only if $\min(F) = \min(f)$.*

Our next structural result is the following lemma. The *support* of a monomial is the set of variables appearing in it with nonzero degree; the size of this set is the *length* of the monomial. A product gh of two polynomials is *m-balanced*, if the minimum length l of one of these polynomials satisfies $m/3 < l \leq 2m/3$. A monomial p *appears m-balanced* in a product gh of two polynomials, if there are monomials $r \in \text{mon}(g)$ and $s \in \text{mon}(h)$ such that rs and p have the same support, and the length l of r satisfies $m/3 < l \leq 2m/3$. Note that here the order of polynomials in their product gh is important: the condition is only on parts of monomials appearing in the first polynomial. In particular, if several monomials appear *m-balanced* in gh , then we know the bounds on the lengths of their parts in *one and the same* of the two polynomials.

Lemma 6. *Let $m \geq 2$, and let f a polynomial of counting complexity $C_{0/1}(f) = s$.*

- (i) *If every monomial of f has length at least m , then $\sup(f)$ is a union of at most s supports of m -balanced products of polynomials.*
- (ii) *There are s products gh of polynomials such that $\sup(gh) \subseteq \sup(f)$, and every monomial of f of length at least m appears m -balanced in at least one of these products.*

Various versions of claim (i) (with degree of or the total number of variables in polynomials used instead of their length) were observed by several authors including Hyafil [9], Jerrum

and Snir [10], Valiant [27], and Raz and Yehudayoff [16]. The advantage of claim (ii) is its wider applicability: the polynomial f itself is allowed to have also short monomials, shorter than m .

Our next results are more *explicit* lower bounds for counting circuits. The r -th degree, $\#_r(A)$, of a family of sets A is the maximum number of sets in A containing a fixed r -element set:

$$\#_r(A) = \max_{|b|=r} |\{a \in A : a \supseteq b\}|.$$

In other words, the intersection of any $\#_r(A)$ sets in A can have at most r elements. Note that

$$|A| = \#_0(A) \geq \#_1(A) \geq \dots \geq \#_r(A) = 1 > 0 = \#_{r+1}(A).$$

where $r = \max\{|a| : a \in A\}$. Also, $A \subseteq B$ implies $\#_r(A) \leq \#_r(B)$. If A is a graph (viewed as a set of edges), then $\#_1(A)$ is the maximum degree of A . In general, $\#_r(A)$ is related with $|A|$ as follows: if A is a family of m -element subsets on $[n]$, then for every $r \leq m$ we have that

$$|A| \binom{m}{r} \leq \#_r(A) \cdot \binom{n}{r}.$$

This can be shown by counting in two ways the number M of pairs (a, b) , where $a \in A$, $|b| = r$ and $a \supseteq b$ holds. By first fixing sets $a \in A$, we get that M is equal to the left-hand side. By fixing sets b , and taking all possible m -element sets a containing b , we get that M is at most the right-hand side.

As we mentioned in the introduction, lower bounds for deciding, and hence, also for counting $\{+, \times\}$ -circuits, can be obtained using lower-bounds arguments for monotone boolean circuits (see, e.g. [12, Chapt. 9] and the literature cited herein), but these are not easy to apply. The reason here lies in a “dual character” of these arguments: in order to obtain a large lower bound of the decision complexity of a polynomial f given by (1), not only the set of monomials of f itself but also that of its “dual” f^* must have some good structural properties. The *dual* f^* of a polynomial

$$f = \sum_{u \subseteq [n]} c_u \prod_{i \in u} x_i \quad \text{is} \quad f^* = \prod_{u: c_u > 0} \sum_{i \in u} x_i.$$

Note that, for every 0-1 input $a = (a_1, \dots, a_n)$, $f(a) = 0$ if and only if $f^*(\bar{a}) \neq 0$, where $\bar{a} = (1 - a_1, \dots, 1 - a_n)$. This holds, because every set in $\text{sup}(f^*)$ intersects every set in $\text{sup}(f)$. More precisely, a general lower bound for deciding $\{+, \times\}$ -circuits is the following.

Theorem 7 ([11]). *Let $f(x_1, \dots, x_n)$ be a polynomial, and $2 \leq r, s \leq n$ be integers. Then for every $A \subseteq \text{sup}(f)$ and $B \subseteq \text{sup}(f^*)$ such that $\#_1(A) \leq |A|/2(s-1)$, we have*

$$D(f) \geq \min \left\{ \frac{|A|}{2(s-1)^r \cdot \#_r(A)}, \frac{|B|}{(r-1)^s \cdot \#_s(B)} \right\}.$$

As shown in [11] (see also [12, Chapt. 9]), this criterion allows to obtain strong (super-polynomial) lower bounds on $D(f)$, and hence, also on $C_{0/1}(f)$ and $C(f)$, for some explicit polynomials. The strength of this criterion lies in the possibility to arbitrarily chose both the parameters r, s as well as sub-families A and B . The weakness, however, lies in its “dual nature” making it not easy to apply: *both* $|A|/\#_r(A)$ *and* $|B|/\#_s(B)$ must be large. It is

usually easy to ensure that $|A|/\#_r(A)$ is large. The problem, however, is with the dual set B , because the set of monomials of the dual polynomial f^* may be rather “messy”, even though the polynomial f itself has a “nice” structure. Say, if $f = \text{Per}_n$, then $|A|/\#_r(A) = n!/(n-r)!$ is large enough already for $A = \text{sup}(f)$. But monomials of f^* correspond then to complements of graphs without perfect matchings, and it is difficult to ensure that $|B|/\#_s(B)$ is also large for some family B of such graphs.

For counting $\{+, \times\}$ -circuits, we have a much more handy lower-bounds criterion, avoiding the need of dual polynomials. By the r -th degree, $\#_r(f)$, of a polynomial f we will mean the r -th degree $\#_r(A)$ of its support $A = \text{sup}(f)$. Thus, if f is multilinear, then $\#_r(f)$ is the maximum number of monomials of f containing a common factor of degree r .

Theorem 8. *Let $f = g + h$ be a polynomial such that every monomial of g has at least $m \geq 2$ variables, and every monomial of h has fewer than $m/3$ variables. Then there is an integer r between $m/3$ and $2m/3$ such that*

$$C_{0/1}(f) \geq \frac{|\text{sup}(g)|}{\#_r(g) \cdot \#_{m-r}(g)}. \quad (3)$$

There is yet another general lower-bounds criterion for monotone arithmetic circuits, due to Gashkov [6], and Gashkov and Sergeev [7]. They call a polynomial f (k, l) -sparse, if

$$\text{mon}(gh) \subseteq \text{mon}(f) \text{ implies } |\text{mon}(g)| \leq k \text{ or } |\text{mon}(h)| \leq l.$$

They proved that $C(f) + 1 \geq |\text{mon}(f)| / \max\{k^3, l^2\}$ holds for every such polynomial. Note that the bound is not trivial, because the fact that $|\text{mon}(g)| \leq k$ or $|\text{mon}(h)| \leq l$ holds does not imply that $|\text{mon}(gh)| \leq kl$ must also hold (because we have an “or”, not “and” here). To obtain a similar lower bound for counting circuits, we will modify their notion of “sparsity”.

Let, as before, $\min(f) \subseteq \text{sup}(f)$ denote the family of all members of $\text{sup}(f)$ which are minimal in the sense that they do not contain any other members of $\text{sup}(f)$. Call a polynomial f (k, l) -free if, for every two polynomials g and h ,

$$\text{sup}(gh) \subseteq \text{sup}(f) \text{ implies } |\min(g)| \leq k \text{ or } |\min(h)| \leq l.$$

The reason to only require $|\min(g)| \leq k$ instead of $|\text{sup}(g)| \leq k$ is that then it is (potentially) easier to show that a given polynomial is (k, l) -free: $|\min(g)|$ can be much smaller than $|\text{sup}(g)|$.

Theorem 9. *Let $1 \leq k \leq l$ be integers. For every (k, l) -free polynomial f , its support $\text{sup}(f)$ is a union of at most $2C_{0/1}(f)$ supports $\text{sup}(gh)$ of products gh of polynomials such that $|\min(gh)| \leq kl^2$. In particular,*

$$C_{0/1}(f) \geq \frac{|\min(f)|}{2kl^2}.$$

Remark 3. The proofs of Theorems 8 and 9 extend to $C_{0/1}(f)$ the arguments used in [6, 7, 13] to lower-bound $C(f)$. The main difficulty with the extension (stipulated by the idempotence axiom $x^2 = x$) is that, unlike the measure $\mu(f) = |\text{mon}(f)|$ (used to lower-bound $C(f)$), the measures $|\text{sup}(f)|$ and $|\min(f)|$ are no more “monotone” in the sense that $\mu(f) \leq \mu(fg)$. To see this, take, for example, $f = x_1 + x_2 + \dots + x_n$ and $g = x_1 x_2 \dots x_n$. Then $|\text{sup}(f)| = n$ but $|\text{sup}(fg)| = 1$.

Remark 4. The proofs of Theorems 8 and 9 are based on the fact (Lemma 5) that, if a $\{+, \times\}$ -circuit counts a polynomial f , then the produced by the circuit polynomial F must satisfy $\text{sup}(F) = \text{sup}(f)$. Thus, these bounds do not extend to monotone *boolean* circuits, where we only have a much weaker property $\min(F) = \min(f)$.

3 Some Applications

Theorem 8 allows us to easily obtain strong lower bounds on $C_{0/1}(f)$ for many polynomials. Let us demonstrate this on some of them. First, associate with every set H of permutations $h : [n] \rightarrow [n]$ the polynomial in n^2 variables $x_{i,j}$:

$$f_H(x) = \sum_{h \in H} \prod_{i=1}^n x_{i,h(i)}.$$

For example, if H consists of all permutations, then f_H is the permanent polynomial Per_n . If H consists of all *cyclic* permutations, then the monomial of f_H correspond to Hamiltonian cycles in K_n .

Corollary 1. *For every set H of permutations of $[n]$, there is an r such that $n/3 < r \leq 2n/3$ and*

$$C_{0/1}(f) \geq \frac{|H| \binom{n}{r}}{n!}.$$

In particular, $C_{0/1}(\text{Per}_n) \geq \binom{n}{r} = 2^{\Omega(n)}$.

Proof. The polynomial f_H has $|H|$ monomials, each specified by a permutation $h \in H$ of $[n]$. If some r variables are fixed, this fixes r values of h . Hence, at most $(n-r)!$ of the permutations can take r pre-described values, implying that $\#_r(f) \leq (n-r)!$. Thus, Theorem 8 gives that $C_{0/1}(f)$ is at least $|H|$ divided by the maximum of $r!(n-r)!$ over all $n/3 < r \leq 2n/3$. \square

In some cases, Theorem 8 allows to even obtain almost optimal bounds. A *partial t -(n, m, λ) design* is a family A of m -element subsets of $\{1, \dots, n\}$ such that any t -element set is contained in at most λ of its members. We can associate with each such design A a multilinear polynomial

$$f_A(x) = \sum_{a \in A} \prod_{i \in a} x_i.$$

Corollary 2. *For every partial t -(n, m, λ) design A with $m/3 \leq t \leq 2m/3$, we have $C_{0/1}(f_A) \geq |A|/\lambda^2$.*

Proof. For all $m/3 \leq r \leq 2m/3$, we have that both r and $m-r$ are at least $m/3$. Thus, the design property implies that both $\#_r(A)$ and $\#_{m-r}(A)$ are at most λ , and the desired lower bound follows directly from Theorem 8. \square

There are many explicit partial designs A with $\lambda \ll \sqrt{|A|}$. For every of them, the counting complexity of the polynomial f_A is almost the same as the number of monomials. To give an example, let n be a prime power, and let A consist of all subsets $a = \{(i, h(i)) : i \in \text{GF}(n)\}$ of the grid $\text{GF}(n) \times \text{GF}(n)$ corresponding to polynomials $h(z)$ of degree at most $d-1$ over $\text{GF}(n)$. Since no two distinct polynomials of degree $< d$ can coincide on d points, we have that no two monomials of f can share d variables in common, A is a partial 1 -($n^2, n, 1$) design, and we obtain $n^d = |A| \leq C_{0/1}(f_A) \leq n^{d+1}$.

Theorem 9 is more difficult to apply than Theorem 8, but it may help for polynomials, on which the latter theorem fails. To demonstrate this, let A be a set of edges of a bipartite point-line incidence graph of a projective plane $PG(2, q)$, introduced by Singer [20]. The nodes on the left-side correspond to $n = q^2 + q + 1$ points x , and those on the left-side to n

lines L , and x and L are adjacent if $x \in L$. Since every line L has $|L| = q + 1$ points, and every point lies in $q + 1$ lines, this is a d -regular graph of degree $d = q + 1 > \sqrt{n}$. Moreover, the graph is $K_{2,2}$ -free (i.e. contains no complete 2×2 subgraphs), because every two points lie in only one line, and every two lines share only one point. For the polynomial

$$f_A(x) = \sum_{uv \in A} x_u x_v,$$

Theorem 8 can only give a trivial lower bound $C_{0/1}(f_A) \geq |A|/d^2 = \Omega(\sqrt{n})$. Indeed, in this case we have $m = 2$, and hence, $r = 1$. But then both $\#_r(f_A)$ and $\#_{m-r}(f_A)$ are equal $d > \sqrt{n}$. On the other hand, it is not difficult to verify that the $K_{2,2}$ -freeness of A implies that the polynomial f_A is (k, l) -free for $k = l = 1$. Thus, Theorem 9 yields an almost optimal lower bound

$$C_{0/1}(f_A) = \Theta(n^{3/2}).$$

As a second example, let us consider the structurally much simpler *triangle polynomial* of $n = 3m^2$ variables with $m^3 = \Theta(n^{3/2})$ monomials:

$$\Delta_n(x, y, z) = \sum_{i,j,k \in [m]} x_{ik} y_{kj} z_{ij}.$$

Schnorr [17] has shown that $C(\Delta_n) = \Theta(n^{3/2})$; this also follows from the lower bound of Gashkov and Sergeev [7] mentioned above, because the polynomial is $(1, 1)$ -sparse: any triangle is uniquely determined by any two of its edges.

Concerning *counting* circuit complexity of $f = \Delta_n$, Theorem 8 can only yield a trivial lower bound $C_{0/1}(f) \geq m^3/m = m^2 = n/3$, because up to m triangles can share a common edge. Still, Theorem 9 (with some more effort) allows us to obtain an almost optimal lower bound.

Corollary 3. *If $f = \Delta_n$, then $C_{0/1}(f) = \Theta(n^{3/2})$.*

Proof. The upper bound $C_{0/1}(f) = O(m^3) = O(n^{3/2})$ is trivial. To prove the lower bound $C_{0/1}(f) = \Omega(m^3)$, we will use Theorem 9. Since $|\text{sup}(f)| = m^3$, it is enough to show that f is $(1, 1)$ -free. To show this, assume that $\text{sup}(gh) \subseteq \text{sup}(f)$ for some polynomials g and h such that $|\text{min}(g)| \geq 2$ and $|\text{min}(h)| \geq 2$. Take any two sets $a_1, a_2 \in \text{min}(g)$, and two sets $b_1, b_2 \in \text{min}(h)$. Then all four unions $a_i \cup b_j$ must be triangles (not just contain a triangle). Moreover, a_1 and a_2 , as well as b_1 and b_2 must be incomparable under inclusion.

Case 1: Some of the sets a_1, a_2, b_1, b_2 forms a triangle T , say $a_1 = T$. Hence, b_1 and b_2 lie in T , and $a_2 \not\subseteq T$ since a_1 and a_2 must be incomparable. Consider the triangles $T_1 = a_2 \cup b_1$ and $T_2 = a_2 \cup b_2$. If $|b_i| \geq 2$ for some $i \in \{1, 2\}$, then $|T_i \cap T| \geq |b_i| \geq 2$, implying that $T_i = T$, and hence, also $a_2 \subseteq T$, a contradiction. So, $b_1 = \{e_1\}$ and $b_2 = \{e_2\}$ for some edges $e_1 \neq e_2$. Since then $|a_2| \geq 2$, the triangles T_1 and T_2 are uniquely determined by a_2 , implying that $T_1 = T_2$ must be the same triangle. But this triangle shares two distinct edges e_1 and e_2 with T , implying that $T_1 = T$, and hence also $a_2 \subseteq T$, a contradiction.

Case 2: None of the sets a_1, a_2, b_1, b_2 forms a triangle. In this case, some of the sets must have exactly two edges, say $a_1 = \{e_1, e_2\}$. Since a triangle is uniquely determined by any two of its edges, we have that both unions $a_1 \cup b_1$ and $a_1 \cup b_2$ must form the same triangle $T = \{e_1, e_2, e_3\}$. The sets b_1 and b_2 must be incomparable, and both of them must contain the “missing” edge e_3 . Since none of these two sets can be a triangle, this implies that $b_1 = \{e_1, e_3\}$

and $b_2 = \{e_2, e_3\}$. These two sets also uniquely determine the same triangle T , implying that $a_2 \cup b_1 = a_2 \cup b_2 = T$. Thus, a_2 must contain both missing edges e_1 and e_2 of T . But this means that a_2 contains the set a_1 , a contradiction with a_1 and a_2 being incomparable. \square

We now turn to the proofs of our main results.

4 Proof of Theorem 2

To show that the gap $C(f_{\text{he}})/C(f)$ can be exponential, consider the following polynomial in $n = m^2 + m$ variables:

$$\text{Per}^*(x, y) = \prod_{i=1}^m \sum_{j=1}^m x_{ij} y_j. \quad (4)$$

The relation to the permanent polynomial Per is that the coefficient of the monomial $y_1 y_2 \cdots y_m$ in $\text{Per}^*(x, y)$ is exactly $\text{Per}_m(x)$.

Now, let $f(x, y)$ be the linearization of $\text{Per}^*(x, y)$. That is, $f(x, y)$ is a multilinear polynomial obtained from $\text{Per}^*(x, y)$ by removing all nonzero exponents from all monomials. Every monomial of f has degree (sum of exponents) between $m + 1$ and $2m$, and the monomials

$$x_{1,j_1} x_{2,j_2} \cdots x_{m,j_m} y_1 y_2 \cdots y_m$$

of degree $2m$ with all j_1, \dots, j_m distinct are exactly the monomials of the polynomial

$$h(x, y) = \text{Per}_m(x) \cdot y_1 y_2 \cdots y_m.$$

Thus, $h = f_{\text{he}}$ is the higher envelope of f . Since $h(x, 1, \dots, 1) = \text{Per}_m(x)$, Theorem 1 yields

$$C_{0/1}(f_{\text{he}}) \geq C_{0/1}(\text{Per}_m) = 2^{\Omega(m)} = 2^{\Omega(\sqrt{n})}.$$

On the other hand, since exponents play no role on 0-1 inputs, we have that $\text{Per}^*(a) = f(a)$ holds for all 0-1 inputs a . Thus, the polynomial f can be counted by the circuit given by the definition (4) of Per^* . This gives the desired upper bound $C_{0/1}(f) = O(m^2) = O(n)$.

To show that the gap $C(g_{\text{le}})/C(g)$ can also be exponential, consider the following polynomial in $n = m^2$ variables x_{ij} given by the formula:

$$\text{Isol}_n(x) = \prod_{i=1}^m \prod_{j=m+1}^{2m} \left(\sum_{k=m+1}^{2m} x_{ik} \right) \left(\sum_{l=1}^m x_{lj} \right). \quad (5)$$

The monomials of this polynomial are obtained as follows. We interpret the variables x_{ij} as edges of a complete bipartite $m \times m$ graph $I \times J$ with parts $I = \{1, \dots, m\}$ and $J = \{m+1, \dots, 2m\}$. To get a monomial of Isol , we take, for each node $i \in I$ exactly one edge x_{ik} incident with i , and take, for each node $j \in J$ exactly one edge x_{lj} incident with j . So, every variable has degree at most 2. Note that on every 0-1 input $a \in \{0, 1\}^n$, $\text{Isol}(a) = 0$ if and only if the graph specified by a has an isolated node.

Let g be the linearization of Isol_n . Every monomial of g has degree between m and $2m$, and the monomials of degree m correspond to perfect matchings. Thus, the lower envelope g_{le} of g is just the permanent polynomial, i.e. $g_{\text{le}} = \text{Per}_m$. By Theorem 1, $C(g_{\text{le}}) = 2^{\Omega(m)}$.

On the other hand, since exponents play no role on 0-1 inputs, we have that $\text{Isol}(a) = g(a)$ holds for all 0-1 inputs a . Thus, the polynomial g can be counted by the circuit given by the definition (5) of Isol . This gives the desired upper bound $C_{0/1}(g) = O(m^3) = O(n^{3/2})$. \square

5 Proof of Theorem 4

Recall that the s - t path polynomial $f = \text{Path}_n$ has one variable $x_{i,j}$ for each edge of a complete undirected graph on $n+2$ nodes $\{s, 1, \dots, n, t\}$. Each monomial of f corresponds to a simple directed path from node s to node t .

The upper bound $D(f) = O(n^3)$ of the *decision* complexity of $f = \text{Path}_n$ follows from the Bellman–Ford dynamic programming algorithm [3, 5]. The circuit is constructed recursively by taking $F_{1,j} = x_{s,j}$ for all $j \in [n] \cup \{t\}$, and using the recursion $F_{l+1,j} = F_{l,j} + \sum_{i \neq j} F_{l,i} \times x_{i,j}$ for $j \in [n] \cup \{t\}$. Monomials of $F_{l,j}$ correspond to walks from node s to node j passing through at most l edges; one edge may be passed more than once, and each pass counts. The output is the polynomial $F = F_{n+1,t}$. Since every s - t walk contains a simple s - t path, and since in deciding $\{+, \times\}$ -circuits we can use the absorption axiom $x + xy = x$, the circuit correctly *decides* Path_n . Thus $D(\text{Path}_n) = O(n^3)$.

Our goal is now to show that every $\{+, \times\}$ -circuit *counting* the s - t path polynomial must have exponential size. We do not have a *direct* proof of this lower bound. Instead, we will derive this result indirectly by using some known reductions and lower bounds.

Say that a $\{+, \times\}$ -circuit *decides* f *with threshold* T , if for every $a \in \{0, 1\}^n$, $F(a) \geq T$ holds precisely when $f(a) \geq 1$. Here, the threshold $T = T(n)$ may depend on the number n of variables, but not on the input. Note that deciding $\{+, \times\}$ -circuits decide with threshold $T = 1$. Let $D_{\text{thr}}(f)$ denote the smallest size of a $\{+, \times\}$ -circuit deciding f with some threshold T .

As defined by Valiant [26], and Skyum and Valiant [21], a polynomial $f(x_1, \dots, x_n)$ is a *monotone projection* of a polynomial $g(y_1, \dots, y_m)$ if there exists an assignment $\sigma : \{y_1, \dots, y_m\} \rightarrow \{x_1, \dots, x_n, 0, 1\}$ such that $f(x_1, \dots, x_n) = g(\sigma(y_1), \dots, \sigma(y_m))$. It is clear that then $D_{\text{thr}}(f) \leq D_{\text{thr}}(g)$.

The r -clique polynomial, $\text{Clique}_{n,r}$, has $\binom{n}{2}$ variables x_e , one for each edge e of K_n , and has one monomial $\prod_{e \in S} x_e$ for every subset $S \subseteq [n]$ of size $|S| = r$. Results of Valiant [25] imply that, for every $1 \leq r \leq n$, $\text{Clique}_{n,r}$ is a monotone projection of the Hamiltonian s - t path polynomial Ham_m for $m = n^{O(1)}$; as noted by Alon and Boppana [1], already $m = 25n^2$ is enough in this case. On the other hand, it is known that, for r about \sqrt{n} , the clique polynomial $f = \text{Clique}_{n,r}$ requires $D_{\text{thr}}(f) \geq 2^{n^{\Omega(1)}}$ [8, 15, 11]; see, e.g. [12, Sect. 9.8] for a simpler proof. (In fact, this result holds for more general circuits where arbitrary monotone real valued functions $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ can be used as gates.) Since $\text{Clique}_{n,r}$ is a monotone projection of Ham_m , we have that

$$D_{\text{thr}}(\text{Ham}_m) \geq D_{\text{thr}}(\text{Clique}_{n,r}) = 2^{n^{\Omega(1)}}.$$

It remains therefore to show that

$$C_{0/1}(\text{Path}_m) \geq D_{\text{thr}}(\text{Ham}_n) \text{ for } m = n^{O(1)}.$$

This can be shown using a standard reduction of Path to Ham . Let $p = (n+1) \log n$. Given an input graph G on $n+2$ nodes $\{s, 1, 2, \dots, n, t\}$, replace each edge (u, v) by a graph on $2p+2$ nodes (u, v and $2p$ new nodes) containing exactly 2^p paths of length $p+1$ between u and v . This way, every s - t path of length l in G gives $(2^p)^l$ s - t paths in the resulting graph G' . This graph has $m = O(pn^2) = O(n^3 \log n)$ nodes.

If G has a Hamiltonian s - t path (of length $n+1$), then the graph G' has at least $T := (2^p)^{n+1}$ s - t paths. If G has no Hamiltonian path, then the longest s - t path has at most n edges, and hence, at most $n-1$ inner nodes. The number of s - t paths of length $\leq n$ is

bounded from above by $n \cdot n^{n-1} = n^n$. So, in this case, G' can have at most $(2^p)^n \cdot n^n = l \cdot n^n / 2^p = T/n$ s - t paths. We have thus shown that every $\{+, \times\}$ -circuit counting Path_m for $m = \Theta(pn^2) = \Theta(n^3 \log n)$ decides Ham_n with threshold $T = (2^p)^{n+1}$. \square

6 Proof of Lemma 5

Let $f(x_1, \dots, x_n)$ be a polynomial in which each variable x_i has degree at most t_i , and let $S_i \subseteq \mathbb{N}$ be arbitrary subsets of sizes $|S_i| \geq t_i + 1$, $i = 1, \dots, n$.

Claim 10 (Folklore). *The polynomial f is uniquely determined by its values on $S_1 \times S_2 \times \dots \times S_n$.*

Proof. Induction on n . For $n = 1$, the claim is simply the assertion that a non-zero polynomial of degree t_1 in one variable can have at most t_1 distinct roots. For the induction step, expand the polynomial f by the variable x_n :

$$f(x_1, \dots, x_n) = \sum_{i=0}^{t_n} f_i(x_1, \dots, x_{n-1}) \cdot x_n^i.$$

For each point $a \in S_1 \times \dots \times S_{n-1}$,

$$f(a, x_n) = \sum_{i=0}^{t_n} f_i(a) \cdot x_n^i$$

is a polynomial of degree at most t_n in one variable, and hence, all its coefficients $f_i(a)$, $i = 0, 1, \dots, t_n$ can be recovered knowing the values $f(a, b)$ for all $b \in S_{n+1}$. Knowing the values $f_i(a)$ for all $a \in S_1 \times \dots \times S_{n-1}$ we can, by the induction hypothesis, recover the polynomials f_i , and hence, the original polynomial f . \square

Now let f and h be two polynomials on the same set of n variables such that $f(a) = h(a)$, and hence, also $\overline{f}(a) = \overline{h}(a)$ holds for all $a \in \{0, 1\}^n$. (Recall that \overline{f} is obtained from f by removing all nonzero exponents.) Since the polynomials \overline{f} and \overline{h} are multilinear, Claim 10 with all $S_i = \{0, 1\}$ yields $\overline{f} = \overline{h}$ (they must coincide as multilinear polynomials), and hence, also $\sup(f) = \sup(h)$ must hold as well.

Let us now prove the second claim of Lemma 5: if f and h are polynomials on the same set of variables, then f and h have the same 0-1 roots if and only if $\min(f) = \min(h)$. The “if” part is trivial, because $f(a) > 0$ happens precisely when $p(a) = 1$ for some monomial $p \in \min(f)$. To prove the “only if” direction, assume that f and h have the same 0-1 roots. Our goal is to show that then $\min(f) = \min(h)$ must hold.

Assume contrariwise that there is a monomial $p \in \min(f)$ whose set of variables X_p belongs to $\min(f)$ but not to $\min(h)$. If $X_q \not\subseteq X_p$ holds for all monomials q of h , then we can set all variables in X_p to 1 and the rest to 0. On the resulting assignment $a = a_p$, we will have $h(a) = 0$ but $f(a) \geq p(a) \geq 1$, a contradiction. Thus, there must be a monomial $q \in \min(h)$ such that $X_q \subset X_p$; the inclusion must be proper, because $X_p \notin \min(h)$. But then on the input a_q , we will have $f(a_q) = 0$ but $h(a_q) \geq q(a_q) \geq 1$, a contradiction again. \square

7 Proof of Lemma 6

We will need the following two simple and well-known facts.

A *subadditive weighting* of a circuit is an assignment of nonnegative numbers (weights) to its gates such that the weight of a gate does not exceed the sum of the weights of its inputs.

Claim 11 (Folklore). *If the output gate gets weight m , and every leaf gets weight at most $2m/3$, then there is a gate of weight larger than $m/3$ and at most $2m/3$.*

Proof. By starting at the output gate, and traversing the circuit by always choosing the input of larger weight, we can find a gate v of weight $> 2m/3$ such that both its inputs u and w have weights at most $2m/3$. By the subadditivity of weighting, at least one of the gates u and w have then weight larger than $(2m/3)/2 = m/3$ and at most $2m/3$. \square

Claim 12 (Folklore). *For every gate u in a $\{+, \times\}$ -circuit producing a polynomial F , the polynomial can be written as $F = PQ + R$, where P is the polynomial produced at u .*

(We use capital letters for polynomials only to stress that they are produced by circuits.)

Proof. If we replace the gate u by a new variable y , the resulting circuit produces a polynomial of the form $yH + R$ for some polynomial H , where R does not contain y (albeit H may contain). It remains to substitute all occurrences of the variable y with the polynomial P produced at the gate u . \square

Proof of Lemma 6(i). For a polynomial f , let $l(f)$ denote the minimum number of variables in a monomial of f . Hence, a product gh of two polynomials is m -balanced, if $m/3 < l(g) \leq 2m/3$. We have to show that, if $l(f) \geq m$ for $m \geq 2$, then $\text{sup}(f)$ is a union of at most $s = C_{0/1}(f)$ supports of m -balanced products of polynomials.

To prove this claim, fix a $\{+, \times\}$ -circuit of size $s = C_{0/1}(f)$ counting f . Define the *weight* of a gate u as $l(P)$, where $P = P_u$ is the polynomial produced at u . Hence, the output gate has weight at least $m \geq 2$, and each input gate has weight 1 (which is $\leq 2m/3$ since $m \geq 2$). Since this weighting is subadditive, Claim 11 gives us a gate u with $m/3 < l(P) \leq 2m/3$. By Claim 12, we can write the produced by our circuit polynomial F as a sum $F = PQ + R$. Hence, $\text{sup}(f) = \text{sup}(F) = \text{sup}(PQ) \cup \text{sup}(R)$, where the product PQ is m -balanced.

The polynomial R is obtained from F by removing some monomials. If R is empty, then we are done. Otherwise, the polynomial R can be produced by a circuit with one gate fewer (gate u is set to constant 0, and disappears). Moreover, $\text{mon}(R) \subseteq \text{mon}(F)$ implies that $l(R) \geq l(F) \geq m$ still holds. So, we can repeat the same argument for the polynomial R , until the empty polynomial R is obtained. \square

Proof of Lemma 6(ii). We will now apply Claim 11 not to the entire circuit but to some its sub-circuits. A *parse-subcircuit* of a circuit F is obtained by setting to 0 one of the two inputs of each sum gate. Such a subcircuit F' can also be defined inductively as follows. The output gate of F is included in F' . If a gate u is already included in F' , and if u is a sum gate, then exactly one of the inputs to u are included in F' . If u is a product gate, then both its inputs are included in F' (see Fig. 2). Note that each parse-subcircuit produces exactly one monomial in a natural way, and that each monomial of the polynomial produced by the entire circuit is produced by at least one parse-subcircuit.

Now let F be a circuit of size $s = C_{0/1}(f)$ counting f , and F be the polynomial produced by F . By Lemma 5, we have that $\text{sup}(f) = \text{sup}(F)$. For every monomial p of F of length at least




Figure 2: A circuit and two its parse sub-circuits producing, respectively, the monomials xz^2 and xy .

m , take some parse-subcircuit F_p producing p , and use Claim 11 to find a gate u in F_p such that the part p' of p produced at u in F_p has length l satisfying $m/3 < l \leq 2m/3$. By Claim 12, we can write the polynomial F as a sum $F = PQ + R$, where P is the polynomial produced at gate u (in the entire circuit). Hence, p appears m -balanced in the product $R_u = PQ$. Since we have at most s products R_u , and since $\text{mon}(R_u) \subseteq \text{mon}(F)$ implies $\text{sup}(R_u) \subseteq \text{sup}(f)$, we are done. \square

8 Proof of Theorems 8 and 3

Define the *join* of two families of sets B and C as the family

$$B * C = \{b \cup c : b \in B, c \in C\}$$

of all possible unions. Note that the support of a product gh of two polynomials is the join of the supports of g and h . Note also that, if no set of B intersects any set of C , then we have an upper bound $|A| \leq \#_{|b|}(A) \cdot \#_{|c|}(A)$ on the size of the join $A = B * C$ holding for all $b \in B$ and $c \in C$. This holds because then $|B| = |B * \{c\}| \leq \#_{|c|}(A)$, and similarly $|C| = |\{b\} * C| \leq \#_{|b|}(A)$. If, however, sets in B and in C intersect, then it may happen that $|B| \gg |B * \{c\}|$. Still, also then we have a reasonable upper bound.

Lemma 13. *Let $B * C$ be a join of two families, and $B * C \subseteq A$. Suppose that every set in $B * C$ has size at least m , and that B or C has a set of size r . Then*

$$|B * C| \leq \#_r(A) \cdot \#_{m-r}(A).$$

Proof. Assume w.l.o.g. that the family B contains a set b of size $|b| = r$, and let $A_b = \{b\} * C \subseteq A$. Associate with every $a \in A_b$ the family

$$C_a = \{c \in C : b \cup c = a\}.$$

These families give a partition of C into $|A_b|$ pairwise disjoint subfamilies. Since all sets in A_b contain the set b of size $|b| = r$, we have that

$$|A_b| \leq \#_r(A).$$

On the other hand, for each $a \in A_b$, all sets in C_a , and hence, also all sets in $B * C_a$ contain the set $a \setminus b$ of size $|a \setminus b| \geq m - r$, implying that

$$|B * C_a| \leq \#_{m-r}(A)$$

holds for all $a \in A_b$. Now, every set $b' \cup c'$ in $B * C$ belongs to $B * C_a$ for $a = b \cup c'$. So,

$$\begin{aligned} |B * C| &\leq \sum_{a \in A_b} |B * C_a| \leq \sum_{a \in A_b} \#_{m-r}(A) \\ &\leq |A_b| \cdot \#_{m-r}(A) \leq \#_r(A) \cdot \#_{m-r}(A). \end{aligned} \quad \square$$

Proof of Theorem 8. Let $f = g + h$ be a polynomial such that $l(g) \geq m \geq 2$, and $l(h) < m/3$; here, as before, $l(f)$ denotes the minimum number of variables in a monomial of f . By Lemma 6(ii), there are $s = C_{0/1}(f)$ products PQ of polynomials such that $\text{sup}(PQ) \subseteq \text{sup}(f)$, and every monomial of g appears m -balanced in at least one of these products.

Claim 14. *If $\text{sup}(PQ) \subseteq \text{sup}(f)$, and if at least one monomial of g appears m -balanced in PQ , then $\text{sup}(PQ) \subseteq \text{sup}(g)$ and $|\text{sup}(PQ)| \leq \#_r(g) \cdot \#_{m-r}(g)$ for some $m/3 < r \leq 2m/3$.*

Proof. To show the inclusion $\text{sup}(PQ) \subseteq \text{sup}(g)$, assume contrariwise that there are $a, a' \in \text{sup}(P)$ and $b, b' \in \text{sup}(Q)$ such that $a \cup b \in \text{sup}(g)$, $m/3 < |a| \leq 2m/3$ but $a' \cup b' \in \text{sup}(h)$. Since $|b'| = l$ for some $l < m/3$, the union $a \cup b'$ has size $l < m/3 < |a \cup b'| \leq 2m/3 + l < m$, and hence, cannot belong to $\text{sup}(f)$, a contradiction with $\text{sup}(PQ) \subseteq \text{sup}(f)$. Thus, $\text{sup}(PQ)$ must lie entirely within $\text{sup}(g)$, as claimed.

To show the upper bound on $|\text{sup}(PQ)|$, let $A = \text{sup}(g)$, $B = \text{sup}(P)$ and $C = \text{sup}(Q)$. Since $l(g) \geq m$, and $\text{sup}(PQ) \subseteq \text{sup}(g)$, we have that every set in $B * C = \text{sup}(PQ)$ has at least m elements. On the other hand, since some monomial of g appears m -balanced in PQ , some set in B must have r elements, for some $m/3 < r \leq 2m/3$. For this r , Lemma 13 yields $|A * B| = |\text{sup}(PQ)| \leq \#_r(A) \cdot \#_{m-r}(A)$, as desired. \square

Thus, every monomial of g belongs to at least one of s products PQ of polynomials such that $|\text{sup}(PQ)| \leq \#_r(g) \cdot \#_{m-r}(g)$ for some $m/3 < r \leq 2m/3$. By taking such an r maximizing $\#_r(g) \cdot \#_{m-r}(g)$, the desired lower bound $s \geq |\text{sup}(g)| / \#_r(g) \cdot \#_{m-r}(g)$ follows. \square

Proof of Theorem 3. Recall that our polynomial f has the form $f = g + h$ with $g = \text{Per}_n$ and $h = \sum_{i,j \in [n]} x_{ij}$. Hence, $l(g) = n$ and $l(h) = 1 < n/3$. By Theorem 8, there is an integer r between $n/3$ and $2n/3$ such that $C_{0/1}(f) \geq |\text{sup}(g)| / \#_r(g) \cdot \#_{m-r}(g) \geq n! / r!(n-r)! = 2^{\Omega(n)}$. On the other hand, on every 0-1 input a , we have that $f(a) = 0$ if and only if $h(a) = 0$, because $g(0, \dots, 0) = 0$. Hence, the circuit h decides f , implying that $D(f) = D(h) \leq n^2$. \square

9 Proof of Theorem 9

By Claim 12, we know that, for every gate u in a given $\{+, \times\}$ -circuit F , the produced by the circuit polynomial F can be written as $F = P_u Q_u + R$, where P_u is the polynomial produced at u , Q_u is the polynomial produced “after” the gate u , and R is the polynomial produced by the circuit after the gate u is replaced with constant 0. For our argument, it will be convenient to introduce the notion of a polynomial Q_e produced after an edge $e = (u, v)$ (see Fig. 3):

$$Q_e = \begin{cases} Q_v & \text{if } v = u + w, \\ Q_v P_w & \text{if } v = u \times w. \end{cases}$$

A set E of edges of F is a *cut*, if every input-output path in F contains an edge in E .




Figure 3: For an edge $e = (u, v)$, the polynomial Q_e produced after e is the polynomial $Q_e = Q_v$ produced after the gate v , if $v = u + w$ is a sum gate, and is $Q_e = Q_v P_w$, if $v = u \times w$ is a product gate, where P_w is the polynomial produced before the gate w .

Claim 15. *If E is a cut, then $\text{mon}(F)$ is a union of $\text{mon}(P_u Q_e)$ over all edges $e = (u, v)$ in E .*

Proof. Take a monomial p of the produced polynomial F , and let F_p be any parse-subcircuit producing p . Since E forms a cut, the graph F_p must contain some edge $e = (u, v) \in E$. Then the monomial p has the form $p = p' p''$ where p' is the monomial produced by the subgraph of F_p rooted in u . Thus p' belongs to the polynomial P_u produced in F before the edge e , and p'' belongs to the polynomial Q_e produced after the edge e . Hence, p belongs to $P_u Q_e$, as desired. \square

Proof of Theorem 9. Let F be a $\{+, \times\}$ -circuit of size $s = C_{0/1}(f)$ counting f , and let F be the polynomial produced by F . By Lemma 5, we know that $\text{sup}(F) = \text{sup}(f)$. Hence, the polynomial F is also (k, l) -free. We first transform the circuit F to a circuit F' as follows. For every product gate $v = u \times w$ in F , one of whose inputs, say u , is *small* in that $|\min(P_u)| \leq l$ holds, we remove the edge (u, v) and replace v by a unary (fanin-1) gate $v = P_u \times w$ of “scalar” multiplication by this fixed (small) polynomial P_u . If both inputs produce small polynomials, then we eliminate only one of them. It is clear that F' produces the same polynomial F . In particular, $\text{sup}(F') = \text{sup}(f)$ holds as well.

Say that an edge $e = (u, v)$ of F' is *light*, if $|\min(P_u Q_e)| \leq kl^2$. To finish the proof of the first claim in Theorem 9, it is enough, by Claim 15, to show that every input-output path in F' must contain at least one light edge.

To show this, take an arbitrary input-output path in F' , and let $e = (u, v)$ be the last edge along this path such that $|\min(P_u)| \leq k$; hence, $|\min(P_v)| > k$. Such an edge must exist because $|\min(x_i)| = 1 \leq k$, and since we can assume that $|\min(F)| > k$ (for otherwise the theorem would trivially hold). Together with $\min(P_v Q_v) \subseteq \min(F)$ and $|\min(P_v)| > k$, the (k, l) -freeness of F implies that

$$|\min(Q_v)| \leq l.$$

If v is a sum gate, then $Q_e = Q_v$, and hence, also $|\min(Q_e)| \leq l$. So, the edge e is light in this case:

$$|\min(P_u Q_e)| \leq |\min(P_u)| \cdot |\min(Q_e)| \leq kl.$$

So, assume that v is a product gate. Let u and w be the inputs to v in the original circuit F . Since $|\min(P_u)| \leq k \leq l$, we have that $|\min(P_w)| \leq l$ must hold as well, for otherwise the edge $e = (u, v)$ could not exist in F' (would be already eliminated when going from F to F'). Hence,

$$|\min(Q_e)| = |\min(P_w Q_v)| \leq l^2.$$

So, the edge e is light also in this case:

$$|\min(P_u Q_e)| \leq |\min(P_u)| \cdot |\min(Q_e)| \leq kl^2.$$

Circuits	$x + x = x$	$x^2 = x$	$x + xy = x$	Property
Computing	—	—	—	$F = f$
Counting	—	✓	—	$\overline{F} = \overline{f}$
Approximating	✓	✓	—	$\sup(F) = \sup(f)$
Tropical	✓	—	✓	$\text{Min}(F) = \text{Min}(f)$
Deciding/Boolean	✓	✓	✓	$\min(F) = \min(f)$

Table 1: Summary of which axioms are allowed (✓) in which kind of $\{+, \times\}$ -circuits. The last column indicates what property the produced by a circuit polynomial F must satisfy; here \overline{f} is the linearization of f obtained by removing all nonzero exponents. Tropical circuits are circuits with $x \oplus y = \min(x, y)$ and $x \otimes y = x + y$ functions as gates. Finally, $\text{Min}(f)$ is the set of all *monomials* of f that contain no other monomial of f as a proper factor. The property $\text{Min}(F) = \text{Min}(f)$ holds only if f is multilinear [10, 13].

Since the total number of edges in F' is at most $2s$, we have thus shown that the support $\sup(F') = \sup(f)$ is a union of at most $2s$ families $\sup(PQ)$ with $|\min(PQ)| \leq kl^2$. Since every minimal set of a union of two families must be minimal in at least one of these families, this implies that $\min(f)$ is contained in (albeit not necessarily equal to) the union of the families $\min(PQ)$. Hence, the desired lower bound $s \geq |\min(f)|/2kl^2$. \square

10 Conclusion and Open Problems

The weakness of monotone arithmetic circuits, i.e. of $\{+, \times\}$ -circuits, *computing* a given polynomial f is stipulated by the fact that the produced by the circuit polynomial F must just (syntactically) coincide with f . In particular, then $\text{mon}(F) = \text{mon}(f)$ must hold. On the other pole are $\{+, \times\}$ -circuits just *deciding* f . These are, in fact, monotone boolean circuits, where the idempotence axiom $x^2 = x$ as well as the absorption axiom $x + xy = x$ can be used, and hence, here we only have a weaker property $\min(F) = \min(f)$. While proving lower bounds in the latter (boolean) model is a relatively difficult task, the severe restriction of the former (arithmetic) model makes this task much easier.

In this paper we considered an intermediate model of *counting* $\{+, \times\}$ -circuits. In this case, it is required that the values of F must coincide with those of f on only 0-1 inputs: on other inputs, the values may be different. Thus, counting circuits are $\{+, \times\}$ -circuits that are allowed to use the idempotence axiom $x^2 = x$ (but not the absorption axiom $x + xy = x$). These circuits have an intermediate structural property that $\sup(F) = \sup(f)$ must hold (Lemma 5). We have shown that counting circuits can be exponentially smaller than computing circuits (Theorem 2), and that deciding circuits can be exponentially smaller than counting circuits (Theorem 3).

A next natural question was whether lower-bounds arguments for the weak (computing) model can be extended to work also for the intermediate (counting) model? We have shown that such an extension is possible for two lower-bounds arguments (Theorems 8–9). In fact, our proofs of these bounds hold for $\{+, \times\}$ -circuits that only “approximate” a given polynomial f in that $\sup(F) = \sup(f)$ holds for the produced by the circuit polynomial F (coefficients play no role in our arguments). Approximating circuits can use both idempotence

axioms $x + x = x$ and $x^2 = x$. (Table 1 summarizes the axioms allowed in various types of circuits.) So, these bounds also hold for $\{\cup, *\}$ -circuits constructing a given family $A \subseteq 2^X$ of subsets of a (fixed) finite set. Inputs are single element sets $\{x\}$ with $x \in X$, and gates are set-theoretic union (\cup) and join ($*$) of families. A special case of Theorem 8 (for $h = 0$) gives that, if every set in A has at least $m \geq 2$ elements, then there is an integer $m/3 < r \leq 2m/3$ such that every $\{\cup, *\}$ -circuit constructing A must have at least $|A|/\#_r(A) \cdot \#_{m-r}(A)$ gates.

A “complementary” in a sense to counting $\{+, \times\}$ -circuits model, also lying between computing $\{+, \times\}$ -circuits and deciding $\{+, \times\}$ -circuits, is that of *tropical* circuits, i.e. $\{\min, +\}$ -circuits. These are $\{+, \times\}$ -circuits, where the sum is interpreted as $\min\{x, y\}$, and the product as $x + y$. Such a circuit *computes* a given polynomial f of n variables, if $\hat{F}(a) = \hat{f}(a)$ holds for all $a \in \mathbb{N}^n$, where \hat{f} is the “tropicalization” of f :

$$f(x) = \sum_{e \in \mathbb{N}^n} c_e \prod_{i=1}^n x_i^{e_i} \quad \text{turns to} \quad \hat{f}(x) = \min_{\substack{e \in \mathbb{N}^n \\ c_e \neq 0}} \sum_{i=1}^n e_i x_i.$$

For example, if $f = xy^2 + 3y^2z^3$, then $\hat{f} = \min\{x + 2y, 2y + 3z\}$. Tropical circuits are important, because many dynamic programming algorithms for minimization problems are just recursively constructed tropical circuits.

The difference from counting $\{+, \times\}$ -circuits is that now the absorption axiom $x + xy = x$ is allowed, but the idempotence axiom $x^2 = x$ is not ($x + x \neq x$ unless $x = 0$). As shown in [10, 13], lower bounds for computing $\{+, \times\}$ -circuits hold also for tropical circuits, as long as the target polynomial f is multilinear: in this case we have that $T(f) \geq C(f_e)$, where $T(f)$ is the minimum size of a tropical circuit computing f . In particular, for polynomials which are multilinear and homogeneous (all monomials have the same number of variables), tropical circuits are no more powerful than computing $\{+, \times\}$ -circuits. Still, for non-homogeneous polynomials, tropical circuits can be exponentially more powerful than even counting $\{+, \times\}$ -circuits. In fact, both gaps $C_{0/1}(f)/T(f)$ and $T(f)/C_{0/1}(f)$ can be exponential, meaning that tropical and counting $\{+, \times\}$ -circuits are incomparable.

Proposition 16. *There are multilinear polynomials f and g of n variables such that both $C_{0/1}(f)/T(f)$ and $T(g)/C_{0/1}(g)$ are $2^{\Omega(\sqrt{n})}$.*

Proof. To show the first gap, consider the permanent polynomial $f = \text{Per}_m + \sum_{i,j=1}^m x_{ij}$ on $n = m^2$ variables. Theorem 3 gives $C_{0/1}(f) = 2^{\Omega(m)}$. But $T(f) \leq m^2 = n$ because f can be computed by a tropical circuit $F = \sum_{i,j} x_{ij}$ whose tropicalization is $\hat{F} = \min_{i,j} (x_{ij})$: since variables cannot take negative values, the minimum will be achieved on a single variable. Thus, $C_{0/1}(f)/T(f) = 2^{\Omega(m)}$.

To show the second gap, take the multilinear polynomial g considered in the proof of Theorem 2. The polynomial g is the linearization of the polynomial Isol_n on $n = m^2$ variables given by (5), and has $C_{0/1}(g) = O(n^{3/2})$. On the other hand, every monomial of g has degree between m and $2m$, and the monomials of degree m correspond to perfect matchings. Thus, the lower envelope g_{le} of g is just the permanent polynomial, i.e. $g_{le} = \text{Per}_m$. Since $C(\text{Per}_m) \geq C_{0/1}(\text{Per}_m) = 2^{\Omega(m)}$ (see Corollary 1) and $T(g) \geq C(g_{le})$, the desired lower bound $T(g) = 2^{\Omega(m)}$ follows. \square

As we mentioned above, $T(f) \geq C(f_{le})$ holds for every multilinear polynomial f . Thus, if the lower envelope f_{le} requires large monotone arithmetic circuits, then the polynomial f itself requires large tropical circuits. This, however, does not hold for polynomials whose

lower envelopes have small $\{+, \times\}$ -circuits. An important example in this respect is the s - t path polynomial $f = \text{Path}_n$. Even though we have $C(f) = 2^{\Omega(n)}$ [10], the lower envelope of f consist of just one variable $x_{s,t}$, implying that $C(f|_e) = 0$. And indeed, the Bellman–Ford algorithm (see Sect. 5) gives $T(f) = O(n^3)$.

Problem 1. Does $T(f) = \Omega(n^3)$ hold for $f = \text{Path}_n$?

This would show that the Bellman–Ford algorithm is optimal, if only Min and Plus operations can be used. It is worth to mention that the optimality of the other prominent dynamic programming algorithm—that of Floyd–Warshall [4, 28] for the all-pairs shortest paths problem—is already known. The corresponding to this problem “polynomial” f_n is actually a set of s - t path polynomials Path_n for all choices of the source and target nodes s and t . Thus, unlike for Path_n , every circuit for f_n must already have $\Omega(n^2)$ distinct output gates. The Ford–Warshall algorithm gives $T(f_n) = O(n^3)$. On the other hand, Kerr [14] has shown that also $T(f_n) = \Omega(n^3)$ holds.

In Sect. 5, we have shown that the monotone counting complexity of Path_n is exponential in n . But, unlike bounds given in Sect. 3, our proof for Path indirect and is based on two rather non-trivial known results: the fact that the clique polynomial Clique requires exponential monotone real circuits, and is a projection of the Hamiltonian s - t path polynomial Ham .

Problem 2. Give a direct proof of $C_{0/1}(f) = 2^{n^{\Omega(1)}}$ for $f = \text{Path}_n$.

Finally, it would be interesting to extend to the case of counting $\{+, \times\}$ -circuits one of the first lower-bounds arguments for computing $\{+, \times\}$ -circuits suggested by Schnorr in [17]. Namely, he proved that $C(f) \geq |\text{mon}(f)| - 1$ holds, if the polynomial f is *separated* in the following sense: for every two monomials $p \neq q$ of f , their product pq does not contain any third monomial $r \notin \{p, q\}$ of f as a factor (see also [13, Sect. 8] for a somewhat simpler proof). This criterion allows to easily prove strong lower bounds for some polynomials. For example, using it, one can easily show that $C(f) \geq \binom{n}{r} - 1$ holds for the r -clique polynomial $f = \text{Clique}_{n,r}$. This polynomial is separated, because the union of no two r -cliques (sets of edges of complete subgraphs of K_n with r nodes) can contain a third r -clique.

Problem 3. Can Schnorr’s argument for $C(f)$ be extended to $C_{0/1}(f)$?

Acknowledgments

I am thankful to Tsuyoshi Ito, Emil Jeřábek, and Igor Sergeev for interesting discussions.

References

- [1] N. Alon and R. Boppana. The monotone circuit complexity of boolean functions. *Combinatorica*, 7(1):1–22, 1987.
- [2] W. Baur and V. Strassen. The complexity of partial derivatives. *Theoret. Comput. Sci.*, 22:317–330, 1983.
- [3] R. Bellman. On a routing problem. *Quarterly of Appl. Math.*, 16:87–90, 1958.
- [4] R.W. Floyd. Algorithm 97, shortest path. *Comm. ACM*, 5:345, 1962.
- [5] L.R. Ford. Network flow theory. Technical Report P-923, The Rand Corp., 1956.
- [6] S.B. Gashkov. On one method of obtaining lower bounds on the monotone complexity of polynomials. *Vestnik MGU, Series 1 Mathematics, Mechanics*, 5:7–13, 1987.
- [7] S.B. Gashkov and I.S. Sergeev. A method for deriving lower bounds for the complexity of monotone arithmetic circuits computing real polynomials. *Math. Sbornik*, 203(10):33–70, 2012 (in Russian). English translation in: *Sbornik: Mathematics*, 203(10) (2012) 1411–1447.

- [8] A. Haken and S.A. Cook. An exponential lower bound for the size of monotone real circuits. *J. Comput. Syst. Sci.*, 58(2):326–335, 1999.
- [9] L. Hyafil. On the parallel evaluation of multivariate polynomials. *SIAM J. Comput.*, 8(2):120–123, 1979.
- [10] M. Jerrum and M. Snir. Some exact complexity results for straight-line computations over semirings. *J. ACM*, 29(3):874–897, 1982.
- [11] S. Jukna. Combinatorics of monotone computations. *Combinatorica*, 9(1):1–21, 1999. Preliminary version: ECCC Report Nr. 26, 1996.
- [12] S. Jukna. *Boolean Function Complexity: Advances and Frontiers*. Springer-Verlag, 2012.
- [13] S. Jukna. Lower bounds for tropical circuits and dynamic programs. *Theory of Comput. Syst.*, 2014. DOI: 10.1007/s00224-014-9574-4.
- [14] L.R. Kerr. *The effect of algebraic structure on the computation complexity of matrix multiplications*. PhD thesis, Cornell Univ., Ithaca, N.Y., 1970.
- [15] P. Pudlák. Lower bounds for resolution and cutting plane proofs and monotone computations. *J. Symb. Log.*, 62(3):981–998, 1997.
- [16] R. Raz and A. Yehudayoff. Multilinear formulas, maximal-partition discrepancy and mixed-sources extractors. *J. Comput. Syst. Sci.*, 77(1):167–190, 2011. Preliminary version in: Proc. of 49th FOCS, 2008.
- [17] C.P. Schnorr. A lower bound on the number of additions in monotone computations. *Theor. Comput. Sci.*, 2(3):305–315, 1976.
- [18] R. Sengupta and H. Venkateswaran. A lower bound for monotone arithmetic circuits computing 0-1 permanent. *Theor. Comput. Sci.*, 209(1–2):389–398, 1998.
- [19] E. Shamir and M. Snir. On the depth complexity of formulas. *Math. Syst. Theory*, 13:301–322, 1980.
- [20] J. Singer. A theorem in finite projective geometry and some applications to number theory. *Trans. Amer. Math. Soc.*, 43(3):377–385, 1938.
- [21] S. Skyum and L. G. Valiant. A complexity theory based on boolean algebra. *J. ACM*, 32(2):484–502, 1985.
- [22] M. Snir. Size-depth trade-offs for monotone arithmetic circuits. *Theor. Comput. Sci.*, 82(1):85–93, 1991.
- [23] V. Strassen. Die Berechnungskomplexität von elementarsymmetrischen Funktionen und von Interpolationskoeffizienten. *Numerische Mathematik*, 20:238–251, 1973.
- [24] P. Tiwari and M. Tompa. A direct version of Shamir and Snir’s lower bounds on monotone circuit depth. *Inf. Process. Lett.*, 49(5):243–248, 1994.
- [25] L. G. Valiant. Completeness classes in algebra. In *Proc. of 11th Annual ACM Symp. on Theory of Computing*, pages 249–261, 1979.
- [26] L.G. Valiant. The complexity of enumeration and reliability problems. *SIAM J. Comput.*, 8(3):410–421, 1979.
- [27] L.G. Valiant. Negation can be exponentially powerful. *Theor. Comput. Sci.*, 12:303–314, 1980.
- [28] S. Warshall. A theorem on boolean matrices. *J. ACM*, 9:11–12, 1962.